On the Relation between Color Image Denoising and Classification

Jiqing Wu, Radu Timofte, *Member, IEEE*, Zhiwu Huang, *Member, IEEE*, and Luc Van Gool, *Member, IEEE*

arXiv:1704.01372v1 [cs.CV] 5 Apr 2017

***Abstract*—Large amount of image denoising literature focuses on single channel images and often experimentally validates the proposed methods on tens of images at most. In this paper, we investigate the interaction between denoising and classification on large scale dataset. Inspired by classification models, we propose a novel deep learning architecture for color (multichannel) image denoising and report on thousands of images from ImageNet dataset as well as commonly used imagery. We study the importance of (sufficient) training data, how semantic class information can be traded for improved denoising results. As a result, our method greatly improves PSNR performance by 0.34 - 0.51 dB on average over state-of-the art methods on large scale dataset. We conclude that it is beneficial to incorporate in classification models. On the other hand, we also study how noise affect classification performance. In the end, we come to a number of interesting conclusions, some being counter-intuitive.**

***Index Terms*—Color denoising, Classification, High rank filters, Sobolev space.**

I. Introduction

Image denoising is a fundamental image restoration task. Typically it is modeled as

$$\mathbf{Y} = \mathbf{X} + \mathbf{N}, \tag{1}$$

where the goal is to recover the clean image $\mathbf{X}$ from the corrupted image $\mathbf{Y}$ affected by the noise $\mathbf{N}$ assumed to be white Gaussian noise.

A large body of work [6], [17], [32], [3], [21], [9], [22] has been published in this area throughout the years. The proposed methods involve all kinds of techniques, varying from the use of self (patch) similarities to random fields formulations. With the recent resurgence of neural networks, researchers also applied shallow and deep convolutional neural networks (CNN) to solve the denoising problem and achieved impressive results [2], [4], [28]. It is worth mentioning that most of the proposed methods in denoising focus on single channel (gray scale) images. The color image denoising is treated as an application of straightforward modified versions of the solutions developed for single channel images and more often the multiple channels are denoised separately by deploying the single channel solutions. A series of Dabov *et al*. works on denoising are a good example of methods developed for single channel denoising and further modified for handling color images. The block-matching and 3D collaborative filtering (BM3D) method for image denoising of Dabov *et al*. [6] was further extended to Color BM3D (CBM3D) by the same authors [5]. Another top color image denoising has been recently proposed by Rajwade *et al*. [20]. They applied higher order singular value decomposition (HOSVD) and reached competitive results at the cost of heavy computational load.

On the other hand, breakthrough results were achieved by the recent and rapid development of the deep CNNs [14], [15] on various vision tasks such as image classification, recognition, segmentation, and scene understanding. Large CNN models [14], [24], [23] with millions of parameters were proven to significantly improve the accuracy over the non-CNN previous solutions. Even deeper CNN architectures [10] have been shown to further improve the performance.

A key contributing factor to the success of CNN, besides the hardware GPU advances, is the introduction of large scale datasets such as ImageNet [7], COCO [16], and Places [31] for classification, detection, segmentation, and retrieval. The availability of millions of train images supports the development of complex and sophisticated models with millions of learned parameters. At the same time the large datasets are reliable and convincing benchmarks for validating the proposed methods. In comparison, the image denoising works still conduct their validation experiments on surprisingly small datasets. For example, in a very recent work [28] 400 images are used for training and 300 images for testing, images selected from [8], [18], while the BM3D method [6] was introduced on dozens of images. In the light of the existing significant data scale gap between the low (*i.e*. image denoising) and high level vision tasks (*i.e*. classification, detection, retrieval) and the data driven advances in the latter tasks we can only ask ourselves: is it really sufficient to validate denoising methods on small datasets?

Commonly the researchers develop image classification algorithms based on the assumption that the images are clean and uncorrupted. However, most of the images are corrupted and contaminated by noise. The sources of corruption are diverse, they can be due to factors such as: suboptimal use of camera sensors and settings, improper environmental conditions at capturing time, post-processing, and image compression artifacts. More importantly, there are evidences [25] showing that CNN models are highly nonlinear and very sensitive to slight perturbation on the image. It leads to the phenomenon that neural network can be easily fooled [19] by manually generated images. Hence, a study of how noisy images and denoising methods impact on the CNN model performance is necessary.

In summary, our paper is an attempt to bridge image denoising and image classification, and our main contributions are as follows:

J. Wu, R. Timofte, Z. Huang and L. Van Gool are with the Department of Information Technology and Electrical Engineering, ETH Zurich, Switzerland, e-mail: {jwu, radu.timofte, zhiwu.huang, vangool}@vision.ee.ethz.ch

1) We propose a novel deep architecture for denoising which incorporates designs used in image classification and largely improve over state-of-the-art methods.
2) We are the first, to the best of our knowledge, to study denoising methods on a scale of million of images.
3) We conduct a thorough investigation on how Gaussian noise affects the classification models and how the semantic information can help for improving the denoising results.

A. Related Work

In the realm of image denoising the self-similarities found in a natural image are widely exploited by state-of-the-art methods such as block matching and 3D collaborative filtering method (BM3D) of Dabov *et al*. [6] and its color version CBM3D [5]. The main idea is to group image patches which are similar in shape and texture. (C)BM3D collaboratively filters the patch groups by shrinkage in a 3D transform domain to produce a sparse representation of the true signal in the group. Later, Rajwade *et al*. [20] applied the same idea and grouped the similar patches from a noisy image into a 3D stack to then compute the high-order singular value decomposition (HOSVD) coefficients of this stack. At last, they inverted the HOSVD transform to obtain the clean image. HOSVD has a high time complexity which renders the method as very slow [20].

Nowadays most of the visual data is actually tensor (*i.e*. color image and video) instead of matrix (*i.e*. grayscale image). Though traditional CNN models with 2D spatial filter were considered to be sufficient and achieved good results, in certain scenarios, high dimensional/rank filter becomes necessary to extract important features from tensor. Ji *et al*. [11] introduced a CNN with 3-dimensional filter (3DCNN) and demonstrated superior performance to the traditional 2D CNN on two action recognition benchmarks. In their 3DCNN model, the output g_i of feature map at (x, y, z) position at i-th CNN layer is computed as follows:

$$g_i^{x,y,z} = \tanh(b_i + \sum_{l=0}^{L-1}\sum_{m=0}^{M-1}\sum_{n=0}^{N-1} w_{i,k}^{l,m,n} g_{i-1}^{(x+l),(y+m),(z+n)}) \tag{2}$$

where the temporal and spatial size of kernel is L and $M \times N$ respectively.

II. Proposed method

We propose a two stage architecture for color image denoising as depicted in Fig. 2. First we convolve the color image with high pass filters to capture the high frequencies and apply our residual CNN model to obtain the intermediate result. In the second stage, we adapt either AlexNet [14] or VGG-16 [23] deep architecture from image classification and stack it on top of our first stage model, along with introducing a novel cost function inspired by Sobolev space equipped norm [1]. As the experiments will show, our proposed method overcomes the regress-to-mean problem and enables the recovery of high frequency details better than other denoising works.

TABLE I
High pass filters for image pre-processing

0	0.5	0
0	-1	0
0	0.5	0

0	0	0
0	-1	0
0	1	0

0	0	0
0	-1	1
0	0	0

0	0	0
0.5	-1	0.5
0	0	0

Fig. 1. Example of pre-processed input with the high pass filters from Table I

A. First Stage

Image Pre-processing In the state-of-the-art (C)BM3D and HOSVD methods the group matching step is undermined by the noise. The higher the noise level is the finding of similar patches gets more difficult. Moreover, many denoising methods have the tendency to filter/regress to mean and to lose the high frequencies, *i.e*. , the output are simply the local mean of the (highly) corrupted image. To address the above issues, we apply the high pass filters (see Tab. I) to each channel of the noisy input. Such operations correspond to the first and second directional derivative *w.r.t*. x and y direction which highlight the high frequencies. The filtered channel responses are concatenated with the noisy channel for each of the image channels and grouped together. The assumption is that the channels of the image are highly correlated and this information can be exploited. See Fig. 1 for RGB channels and corresponding filter responses of 'Lena'.

Learning correlation by high rank filters As mentioned before, color image is actually tensor data. If for grayscale image it is sufficient to apply spatial filters to extract useful features, for color image the inter-channel correlation is key to the denoising quality. Our later experiments confirms that denoising each image channel independently leads to much poorer performance. Furthermore, after our pre-processing step we gain extra image gradients information, which means our input currently has width, height, channel and gradient dimension. Thus, it does not suffice to use spatial filters in our case. In order to utilize the correlation among RGB channels as well as gradients, we are motivated to apply high rank filters to convolve high rank input. This computation has a nice interpretation based on tensor calculus, which we give a brief introduction here.

Recall 1: A tensor T of rank (r, s) is a multilinear map

$$T : \mathbf{V}^* \times \mathbf{V}^* \times \cdots \times \mathbf{V}^* \times \mathbf{V} \times \mathbf{V} \times \cdots \times \mathbf{V} \to \mathbb{R} \tag{3}$$

with s copies of vector space $\mathbf{V}$ and r copies of its dual $\mathbf{V}^*$. We also consider the set of all tensors of rank (r, s) is a vector space over $\mathbb{R}$, equipped with the pointwise addition and scalar multiplication. This vector space is denoted by $\mathbf{V} \otimes \mathbf{V} \otimes \cdots \otimes$

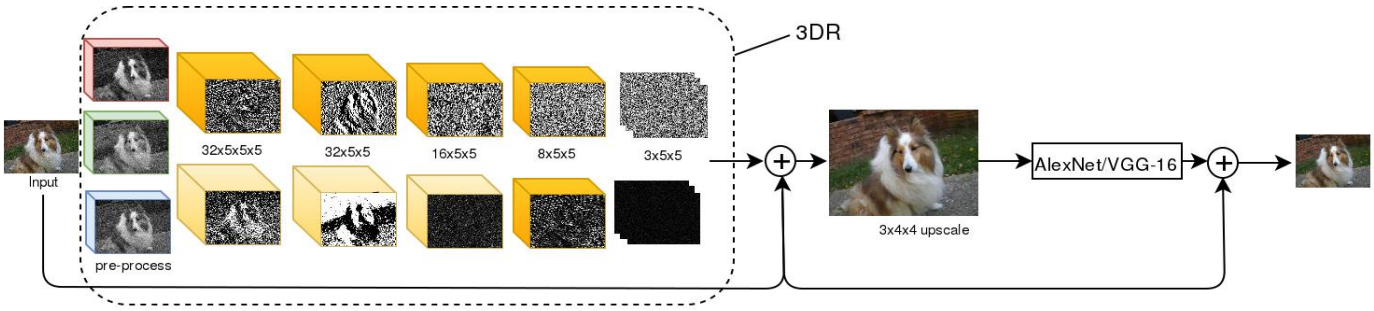

Fig. 2. Outline of proposed model.

$\mathbf{V} \otimes \mathbf{V}^* \otimes \mathbf{V}^* \otimes \cdots \otimes \mathbf{V}^*$ with r copies of $\mathbf{V}$ and s copies of $\mathbf{V}^*$. Moreover, we use Einstein's summation convention $a_r b^r = \sum_{i=0}^{n} a_i b^i$. Suppose T has rank $(2,1)$, then we define

$$T_i^{jk} = T(\mathbf{e}_j, \mathbf{e}_k, \mathbf{e}^i), \tag{4}$$

and we can express $T = T_i^{jk} \mathbf{e}_j \otimes \mathbf{e}_k \otimes \mathbf{e}^i$, where $\mathbf{e}_n, \mathbf{e}^n$ are the basis *w.r.t.* $\mathbf{V}, \mathbf{V}^*$.

Remark 1: (Contraction) New tensors of rank $(r-1, s-1)$ can be produced by summing over one upper and low index of T and obtain $T_{b_1 \cdots b_{j-1} a b_{j+1} \cdots b_s}^{a_1 \cdots a_{i-1} a a_{i+1} \cdots a_r}$.
Now the convolution of n-dimensional filter (rank $(0,n)$) can be considered as the contraction of tensor product between kernel w and small patch g sliced from the input tensor, that is,

$$g^{a_1 \cdots a_n} w_{b_1 \cdots b_n} \mathbf{e}_{a_1} \otimes \cdots \otimes \mathbf{e}_{a_n} \mathbf{e}^{b_1} \otimes \cdots \otimes \mathbf{e}^{b_n} \mapsto g^{a_1 \cdots a_n} w_{a_1 \cdots a_n}. \tag{5}$$

Both width and depth of the CNN matter [29], [10] and considering the trade-off between time complexity and memory usage, our first stage (3DR) is designed for ensemble learning and has two 5-layers with same architecture (See Fig. 2). To incorporate the preprocessed image tensors in our two CNN models, we use the n-dimensional filters to convolve the input, so that we can intensively explore the high frequency information which are greatly contaminated by noise. For efficiency we recommend 3D filters for the first layer of both two CNN models.

We chose tanh instead of ReLU as the activation function mainly for two reasons. First, the negative update is necessary for computing the residual in our denoising task, while ReLU simply ignores the negative value. Second, tanh acts like a normalization of feature map output by bringing the value into the interval $[-1, 1]$ and the extreme outputs will not occur during the training.

Recently, learning residuals has been proved to be a simple and useful trick in previous works [10], [12]. Hence, instead of predicting denoised image, we estimate the image residual. In the end, we simply average the output residuals from the two 5-layer CNN networks $\mathcal{R}_{1,2}$, then obtain the intermediate denoised image $\hat{\mathbf{X}}$ by adding averaged residuals $\mathcal{R}_{1,2}(\mathbf{Y})$ and noisy image$\mathbf{Y}$,

$$\hat{\mathbf{X}} = \lambda_1 \mathcal{R}_1(\mathbf{Y}) + (1 - \lambda_1)\mathcal{R}_2(\mathbf{Y}) + \mathbf{Y}, \tag{6}$$

where λ_1 is fixed to 0.5 in all our experiments. For better summarizing the property of our first stage, we call it 3D residual learning stage (3DR).

B. Second Stage

Collaboration with classification architecture Inspired by [10], we fully take advantage of sophisticated CNN architectures proposed for image classification and adapt them on top of our 3DR model as a second stage. We adapt AlexNet [14] and VGG-16 [23] which are two widely used CNN architectures. We intend to have an end-to-end model and an output denoised image with the same size as the input image. In AlexNet/VGG-16 the stride set in the pooling and convolution layer causes image size reduction. Therefore, we upscale the first stage denoised image with a deconvolution layer and keep only one max pooling layer for AlexNet/VGG-16, such that the size of output image keeps the same. Due to memory constraints, we use only a part of the VGG-16 model up to conv3-256 layer (see [23]). Additionally, we replace the fully connected and softmax layer for both AlexNet and VGG-16 with a novel loss layer matching the target of recovering the high frequencies.

Mixed partial derivative loss Generally, differentiability is an important criteria for studying function space, especially for differential equations, which is the motivation of introducing Sobolev space [1]. Roughly speaking, Sobolev space is a vector space equipped with L_p norm *w.r.t.* function itself and its weak derivatives. Motivated by the norms equipped for the Sobolev space, we propose the so-called mixed partial derivative norm for our loss function.

Recall 2: Let $\alpha = (\alpha_1, \cdots, \alpha_n)$ be the multi-index with $|\alpha| := \alpha_1 + \cdots + \alpha_n$, and f is n-times differentiable function, then the mixed partial derivative is defined as follows

$$D^\alpha f = \frac{\partial^{|\alpha|} f}{\partial x_1^{\alpha_1} \cdots \partial x_n^{\alpha_n}}. \tag{7}$$

Next we introduce our derivatives norm

$$\|f\|_k = \sum_{|\alpha| \le k} \|D^\alpha f\|_2, \tag{8}$$

where 2 indicates the Euclidean norm. Given that we mainly deal with images and obtain its corresponding derivatives by discrete filters, Eq. 8 can be converted to the following

$$\|I\|_k = \sum_{|\alpha| \le k} \|[D]^\alpha * I\|_2, \tag{9}$$

where $[D]$ indicates discrete derivative filter and I is the image. The above formulation is consistent with the preprocessing step by high pass filters (Tab. I), which are exactly the discrete first- and second derivative operators along x and y direction.

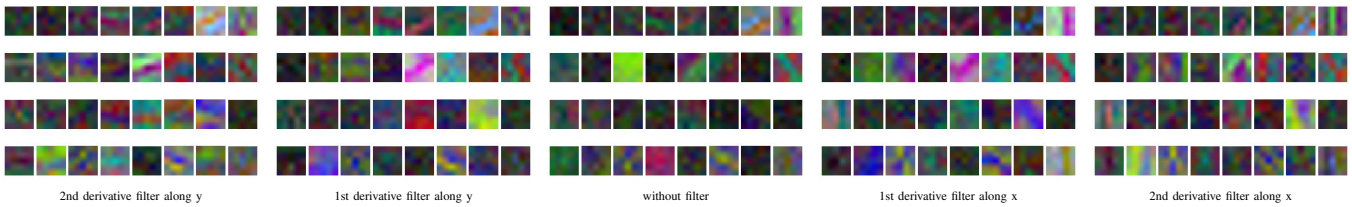

Fig. 3. Visualization of the first layer of 3DR in our model Best zoom on screen.

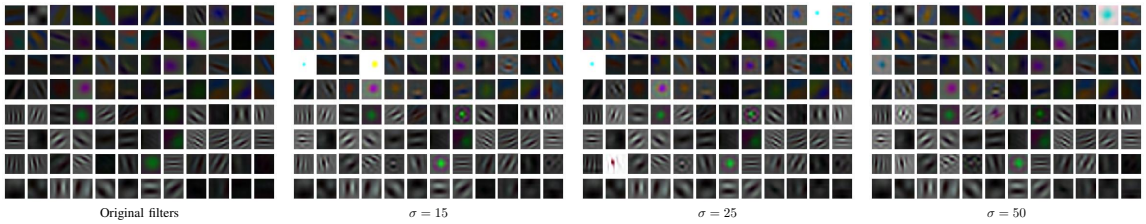

Fig. 4. Visualization of the first layer of AlexNet. Best zoom on screen.

By introducing the mixed partial derivative norm as our loss function, we impose strict constraints on our model so that it decreases the 'derivative distances' between denoised output and clean image and keeps the high frequency details. In our experiments, we set $k = 2$ and ignore the second derivative regarding x and y. Since Peak Signal to Noise Ratio (PSNR) is the standard quantitative quality measure of the denoising results, we combine PSNR with the mixed partial derivative norm to obtain the loss function:

$$L(\mathbf{\Theta}) = \sum_{\alpha \neq (1,1)}^{|\alpha| \leq 2} -20\log_{10}\lambda \| [D]^{\alpha} * (\mathcal{R}(\hat{\mathbf{X}}; \mathbf{\Theta}) + \hat{\mathbf{X}} - \mathbf{X}) \|_2, \tag{10}$$

where $\lambda = \frac{1}{2N}$ and N is the image pixel numbers.

C. Hyperparameters setting

In order to demonstrate the potential of our model, we optimize the model hyperparameters by validation. To this end, we conduct multiple experiments on the training data provided by ImageNet. Meanwhile, we randomly select a small subset of validation data from ImageNet. According to the validation results, we determine important hyperparameters key to the denoising results.

Filter size To figure out the spatial size of our 3DR filters we learn 3DR and report the performance with 3×3, 5×5 and 7×7 filter sizes. As shown in the top left plot of Fig. 5, our favorite spatial size is 5×5. As shown in Fig. 5 5×5 and 7×7 are comparable and both outperform the 3×3 configuration. We set our 3DR filters to 5×5 size as it is less time to compute than 7×7.

Residual vs. Image Residual learning is a useful trick to overcome the accuracy saturation problem, demonstrating its strength in various tasks including image super-resolution [12]. Our experiments on denoising task substantiate the effectiveness of residual learning, Fig. 5 presents an obvious PSNR gap between working on residuals and when using the image itself.

AlexNet and VGG-16 We verify the denoising performance obtained by AlexNet and VGG-16. During the training of second stage, we simply follow their original parameter setup and freeze the 3RD weights. We also compare AlexNet with finetuning of the pretrained ImageNet model (for classification) and AlexNet with training from scratch. Experimental results in Fig. 5(b) show that both of them achieve competitive results and boost the results with a large margin over the first stage. Using the pretrained weights with finetuning seem to lead to a more stable training. Therefore, by default we apply the finetuning strategy for both AlexNet and VGG-16.

Training For training we use Adam [13] instead of the commonly used stochastic gradient descent (SGD). With Adam we achieve a rapid loss decrease in the first several hundreds iterations of the first 3DR stage and further accuracy improvement during the second stage. With SGD we achieve significantly poorer results than with Adam. By cross validation we determine the learning rate to be $0.005, 0.001$ for the first and second stage, and momentum to be 0.9 throughout the stage. The minibatch size is set to 10 and the number of iterations are set to be $200000, 90000$ accordingly for the first and second stage. 100 training image brings the performance to 32.23dB (see Fig. 5(c), slightly better than the first stage result and indicates a clear overfitting and saturation in performance. Increasing the train pool of images to 1000 brings a consistent $\sim 0.4dB$ improvement close to the maximum performance achieved when using 1 million train images. We conclude that for denoising it is necessary to use large sets for training, in the range of thousands. Though the improvement margin is relative small between using 1000 and 1000000 train images, it is no harm to train on million of images for increased robustness. While for all the experimental results we kept the models trained with 290000 iterations (due to lack of time and computational resources) we point out that the performance of our models is still improving with the extra number of iterations ($\sim 0.1dB$ in Fig. 5(d) for 500000 iterations over the 290000 iterations operating point).

D. Architecture Design Choices

In this section, we report a couple of experimental results which support the architecture decisions we made.

Architecture Redundancy Firstly, for 3DR we picked up a design redundancy of two 5-layer networks for which

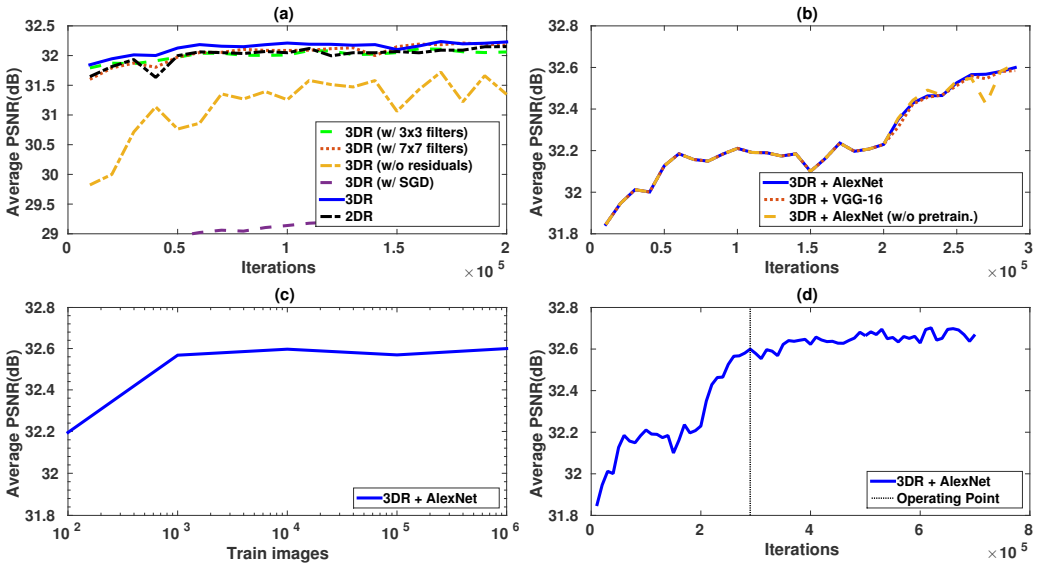

Fig. 5. The denoising performances *w.r.t.* different hyperparameter configuration.

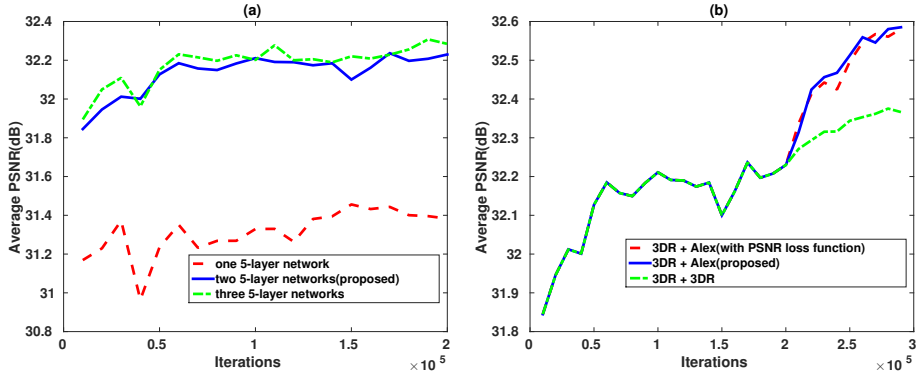

Fig. 6. The denoising performances vs. architecture design on the same validation dataset and noise level $\sigma = 15$. (a) shows how the number of redundant 5-layer network architecture design affects average PSNR, we compare the cases of one 5-layer network, the combination of two and three 5-layer network; (b) shows how loss function and depth of 3DR affects average PSNR, we compare the case of using PSNR and mixed partial derivative loss, and cascading 3DR twice.

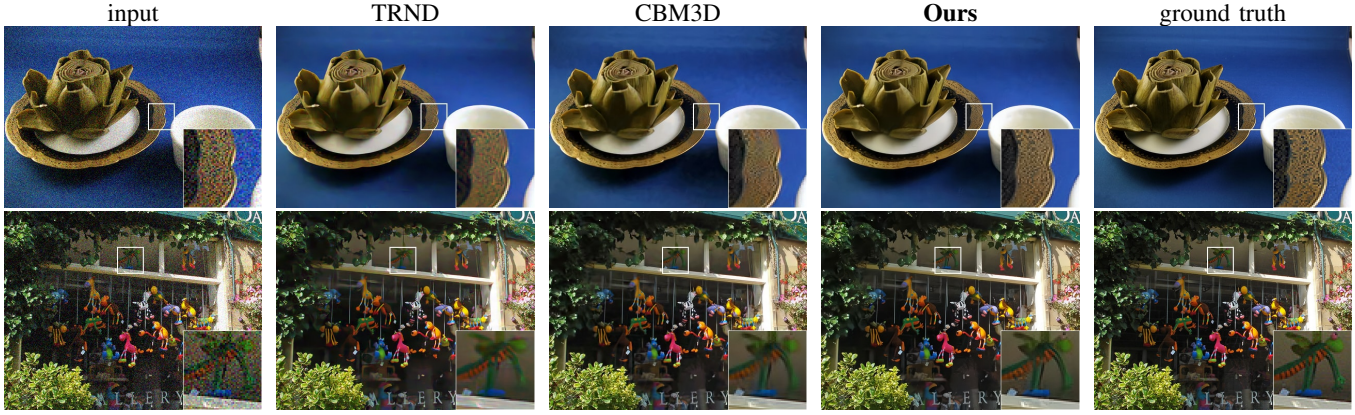

Fig. 7. Visual assessment of non-blind denoising results for $\sigma = 25$ on ImageNet validation. Best zoom on screen.

we then average their residuals. In Fig. 6(a) we report the effect on the PSNR performance of 3DR with 1, 2, and 3 such 5-layer networks on the same validation data(*i.e.* small subset of validation images in ImageNet) and noise level(*i.e.* $\sigma = 15$) mentioned in our paper. The three 5-layer networks architecture of 3DR leads to the best performance for most cases of training iterations. The one 5-layer architecture is much worse than the two 5-layer.

If we apply three 5-layer, we can still slightly improve the PSNR over the two 5-layer design, however the performance gain begins to saturate at the cost of extra computational time. Therefore the two 5-layer networks are the default setting for

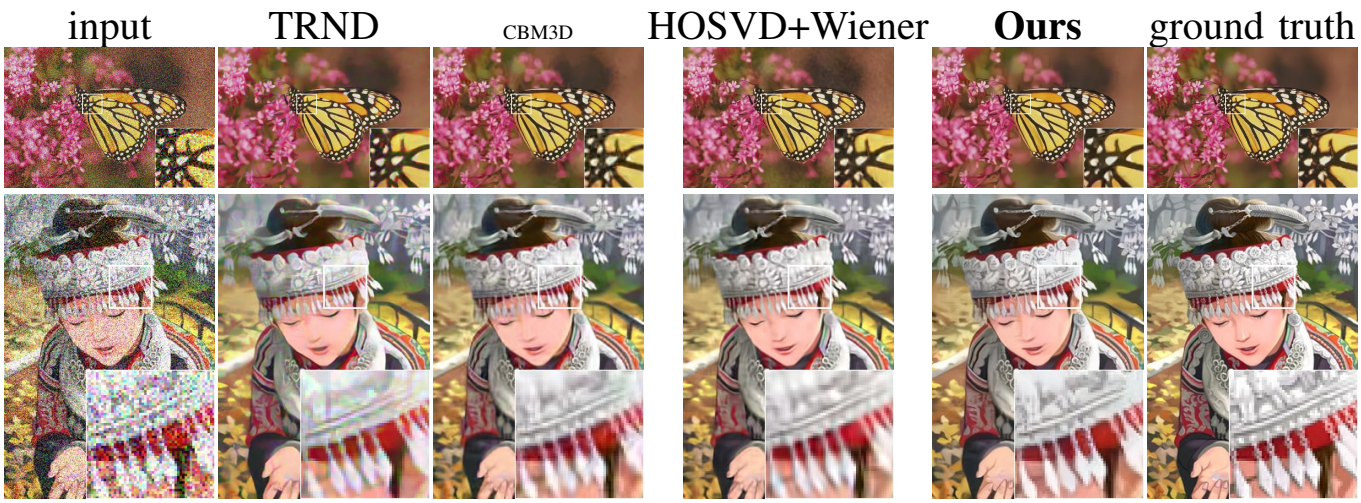

Fig. 8. Visual assessment of non-blind denoising results for $\sigma = 50$ on Set 14. Best zoom on screen.

Fig. 9. Visual assessment of non-blind denoising results for $\sigma = 50$ on ImageNet. Best zoom on screen.

our 3DR model and if sufficient time and memory budget is available then by increasing the number of 5-layer networks in the redundant design of 3DR some improvements are clearly possible.

Architecture Components As analyzed above we propose a 3DR component which can be trained separately and embeds the 3D filtering and joint processing of the color image channels. On top of this standalone component we can cascade either other 3DR components or components derived from published architectures such as AlexNet and VGG-16 proved successful for classification tasks. In Fig. 6(b) we report the results on the same dataset and noise level as Fig. 6(a) when we stack on top of our first 3DR component either: i) another 3DR component or ii) the adapted AlexNet architecture as described in the paper. In both cases we benefit from deepening our network design. However, AlexNet is deeper than 3DR and with fewer training iterations leads to significantly better PSNR performance. Thus, it is helpful to introduce a deep architecture (from classification) as our second stage. Of course, with enough computation resources, we can still gain PSNR improvement by cascading 3DR multiple times or stacking more layers in 3DR.

Loss Function When it comes to the loss function, we prefer our mixed partial derivative (Eq. 10) loss to a PSNR loss. In Fig. 6(b) we compare the PSNR performance of our proposed architecture (3DR+AlexNet) when using the mixed partial derivative loss and when using the PSNR on the second stage/component – AlexNet. For the first 3DR stage we always use PSNR loss. We consider the two stages of our model to be coarse to fine, thus during the coarse stage we use the normal PSNR loss. A mild improvement is achieved when using the mixed partial derivative loss over PSNR loss shown

in Fig. 6(b).

III. Experiments

A. Image denoising

Our experiments are mainly conducted on the large ImageNet dataset. We train our models with the complete ImageNet training data containing more than 1 million images. Due to the GPU memory limitation and for training efficiency, we randomly crop the training image with size 180×180 and set minibatch to be 10. For testing we have 1000 images with original size, collected from the validation images of ImageNet, which have not been used for the hyperparameter validation. Complementary to the ImageNet test images we also test our methods on traditional denoising and reconstruction benchmarks: Kodak and Set14 [30], [27], which contains 24 and 14 images, resp. Set14 contains classical images such as Lena, Barbara, Butterfly.

TABLE II
THE AVERAGE PSNR AND SSIM PERFORMANCES FOR KODAK DATASET. THE BEST RESULTS ARE IN BOLD.

Methods	σ = 15		25		50		90		130	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
TRND	32.48	0.9420	30.10	0.9085	27.23	0.8464	-	-	-	-
CBM3D	34.43	0.9621	31.83	0.9377	28.63	0.8865	26.27	0.8288	23.48	0.6929
HOSVD	33.66	0.9540	31.12	0.9248	27.97	0.8674	-	-	-	-
HOSVD+wiener	34.18	0.9594	31.64	0.9332	28.38	0.8781	-	-	-	-
3DR+VGG-16(b)	33.99	0.9587	31.84	0.9373	28.64	0.8824	19.87	0.4724	-	-
3DR+AlexNet(b)	33.98	0.9587	31.87	0.9380	28.72	0.8856	19.40	0.4975	-	-
3DR+VGG-16	34.66	0.9635	32.11	0.9403	28.88	0.8891	26.37	0.8243	24.85	0.7723
3DR+AlexNet	34.69	0.9636	32.16	0.9409	28.98	0.8906	26.46	0.8282	25.00	0.7789
3DR+VGG-16(r)	34.75	0.9641	32.20	0.9413	28.98	0.8912	26.48	0.8282	24.97	0.7772
3DR+AlexNet(r)	**34.79**	**0.9643**	**32.27**	**0.9423**	**29.06**	**0.8932**	**26.61**	**0.8335**	**25.15**	**0.7860**

TABLE III
THE AVERAGE PSNR AND SSIM PERFORMANCES FOR SET14 DATASET. THE BEST RESULTS ARE IN BOLD.

Methods	σ = 15		25		50		90		130	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
TRND	31.88	0.9450	29.54	0.9125	26.49	0.8472	-	-	-	-
CBM3D	33.18	0.9610	30.84	0.9375	27.83	0.8875	25.31	0.8267	22.67	0.7211
HOSVD	32.77	0.9553	30.51	0.9285	27.34	0.8706	-	-	-	-
HOSVD+wiener	33.15	0.9596	30.91	0.9360	27.77	0.8834	-	-	-	-
3DR+VGG-16(b)	32.34	0.9539	30.69	0.9352	27.76	0.8864	19.67	0.6079	-	-
3DR+AlexNet(b)	32.33	0.9538	30.71	0.9356	27.85	0.8887	19.32	0.5929	-	-
3DR+VGG-16	33.38	0.9624	31.05	0.9394	27.97	0.8901	25.29	0.8248	23.68	0.7739
3DR+AlexNet	33.41	0.9625	31.11	0.9400	28.01	0.8904	25.38	0.8271	23.81	0.7776
3DR+VGG-16(r)	33.48	0.9631	31.18	0.9407	28.10	0.8923	25.46	0.8292	23.86	0.7798
3DR+AlexNet(r)	**33.53**	**0.9634**	**31.26**	**0.9416**	**28.17**	**0.8932**	**25.58**	**0.8330**	**24.01**	**0.7852**

TABLE IV
THE AVERAGE PSNR AND SSIM PERFORMANCES FOR 1000 IMAGES FROM VALIDATION DATA OF IMAGENET. THE BEST RESULTS ARE IN BOLD.

Methods	σ = 15		25		50		90		130	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
TRND	31.54	0.9328	29.11	0.8951	26.12	0.8231	-	-	-	-
CBM3D	32.98	0.9504	30.47	0.9215	27.30	0.8604	24.94	0.7930	22.54	0.6672
3DR+VGG-16	33.29	0.9528	30.86	0.9262	27.76	0.8710	25.27	0.8017	23.82	0.7499
3DR+AlexNet	**33.32**	**0.9530**	**30.92**	**0.9269**	**27.81**	**0.8713**	**25.35**	**0.8041**	**23.94**	**0.7539**

Non-blind denoising Most typical setup is when the noise is Gaussian with known standard deviation σ. Given a wide range of $\sigma = 15, 25, 50, 90, 130$ we compare our methods with CBM3D, recently proposed HOSVD [20] [1] and its variant HOSVD with wiener filter (HOSVD+Wiener), as well as one state-of-the-art grayscale image denoising method TRND [4] applied independently on each image channel. The performance on Set14, Kodak and ImageNet datasets is reported in Tab. III, II, and IV. Our proposed methods reach the best performances for all testing datasets. Fig. 5(a) reports 3DR performance and Fig. 5(b) shows further gains with AlexNet/VGG-16. Generally the full 3DR+AlexNet/VGG-16 method are 0.2-0.4dB higher than 3DR. On ImageNet our 3DR+AlexNet goes from an improvement of 0.34dB PSNR over CM3D for small levels of noise ($\sigma = 15$) to 0.51dB for medium levels of noise ($\sigma = 50$) and to 1.5dB for high noise ($\sigma = 130$). The larger the noise (σ) is, the larger the performance gap gets. Thus, our methods cope well with high levels of noise while for CBM3D it is getting more difficult to correctly group similar patches. 3DR+AlexNet performs sligtly better than 3DR+VGG-16. Our models perform the best also when using the SSIM measure, which in our experiments correlates well with PSNR. On Kodak and Set14 the denoising results and the relative performance are consistent with the ones on ImageNet. Considering also that our methods were trained on ImageNet train dataset which has a different distribution than Kodak or Set14 we can conclude that our methods, 3DR+AlexNet/VGG-16, do not overfit ImageNet images and generalizes well for color images outside this large dataset. Tab. II and III also shows that our methods are better than the recent HOSVD methods, which cost hours to process a single image, and that using the inter-channel correlations for color image denoising is a must, otherwise top state-of-the-art single channel denoising methods such as TRND provide poor reference performance. In addition, we use enhanced prediction trick [26] and rotate the noisy images by 0, 90, 180, 270 degrees, process them and then rotate back and average the denoised outputs. In this way we achieve significant performance gains for both our methods 3DR+AlexNet(e) and 3DR+VGG-16(e) at the cost of increased running time. For instance, on Set14, 3DR+AlexNet(e) gains 0.16dB over 3DR+AlexNet at $\sigma = 50$ and 0.2dB at $\sigma = 130$. Visual results also confirm the quantitative improvements achieved by our methods (See Fig. 7, 8, 9)

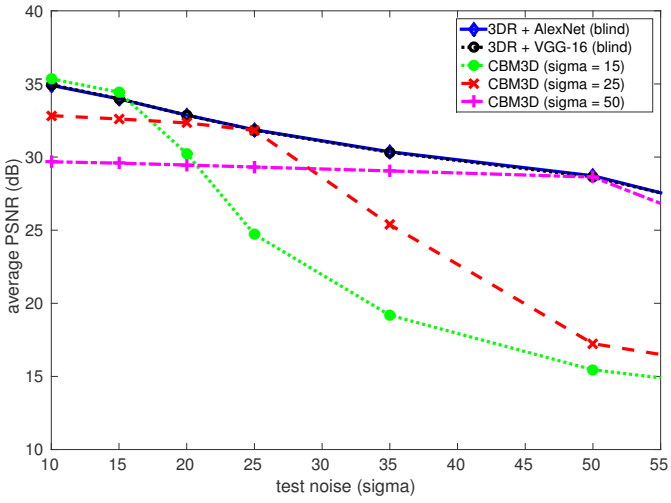

Fig. 10. Blind denoising comparison of our 3DR+AlexNet and 3DR+VGG-16 blindly trained methods and CBM3D with fixed noise levels $\sigma \in \{15, 25, 50\}$

Blind denoising In practice, the noise level is seldom known and can vary with the source of noise and time. Therefore,

[1] Unfortunately, due to high time complexity of HOSVD we did not conduct its experiments on ImageNet.

Fig. 11. Real image denoising, the best CBM3D result is shown. Best zoom on screen.

the methods are able to cope 'blindly' with unknown levels of noise are very important for real applications. As shown in Fig. 10 for the reference CBM3D method whenever there is a mismatch between the level of noise in the test image and the one the method is set for the performance significantly degrades in comparison with the same method with 'knowledge' of the noise level at test.

In order to test the robustness of our proposed denoising methods (3DR+AlexNet and 3DR+VGG-16) to blind denoising conditions we train using the same settings as before, for the non-blind denoising case. The only difference is that the noise level randomly changes from one train image sample to another during the training. In this way the deep model learns to denoise Gaussian noise 'blindly'. In Fig. 10 we compare the performance of our blind methods with that of the CBM3D method with knowledge of 3 levels of noise, $\sigma \in \{15, 25, 50\}$. The test images are the 24 Kodak images. Our blind methods perform similarly regardless the level of noise and are effective envelopes for the performance achieved using CBM3D with various noise settings. Note that only for the low levels of noise ($\sigma < 15$) our blind methods perform poorer than the non-blind CBM3D. At the same time, as expected, our blind models (3DR+VGG-16(b), 3DR+AlexNet(b)) achieve a denoising performance below the one of our non-blind models as numerically shown in Tab. III and II.

We also present here a visual result for real image denoising Fig. 11, which means we have no clue about the noise pattern. The image was taken under dim light conditions with truly poor image quality. Visual result shows that our method still provides better image quality compared to the best CBM3D result.

Running time We use the GeForce GTX TITAN X GPU card to train and test our models. It takes about one day for the 200,000 iterations of the 3DR training phase. Another 1.5 days are necessary for the following 90,000 iterations to train a whole model (3DR+AlexNet or 3DR+VGG-16). At test, an image with size 500×375 is processed in about 60 seconds, including CPU and GPU memory transfer time. In comparison, HOSVD costs about 4 hours to denoise the same image on CPU.

Visualization We visualize the first layer filters of 3DR and finetuned AlexNet (See Fig. 3 and 4) to gain some insights of proposed model. Filters on noisy image itself (middle plot of Fig. 3) shows that color is useful for denoising, there exist green and pink filter among other edge filters. Moreover, the filters obtained by finetuning ImageNet (Fig. 4) seems quite interesting. We can easily notice that most of the finetuned

Example	Category	general training	Category training
	Labrador	33.38	**33.44**
	Ski	33.41	**33.46**
	Tram	32.13	**32.30**
	Broccoli	31.62	**31.67**
	Dining table	33.45	**33.47**
	Airliner	34.49	**34.62**

Fig. 12. PSNR performances achieved by generic - and semantic label finetuning.

class	Airliner			Dining table		
image				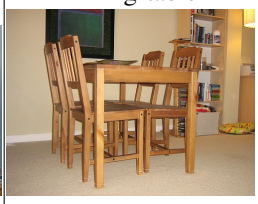		
mask						
generic	36.60	32.53	38.36	34.21	33.62	34.77
semantic	36.66	32.62	38.37	34.24	33.67	34.79
improvement	+0.06	+0.09	+0.01	+0.03	+0.07	+0.02

Fig. 13. Denoising PSNR(dB) results for two examples when 3DR + AlexNet is trained using generic settings (no semantic knowledge) and when is trained for each semantic class individually. $\sigma = 15$

filters (2-4 plot of Fig. 4)share lots in common with the original classification filters of AlexNet. It suggests that filters trained for high level task can also help the low level task such as denoising. There are only few filters with white background which are completely different from the original ones, we believe those filters are particularly adapted to detect Gaussian noise.

B. Image classification

In this section, we study how denoising methods affect classification performance. We firstly use various denoising methods to obtain clean image then classify the image by AlexNet and VGG-16. The results are reported on Tab. V

TABLE V
COMPARISON WITH DIFFERENT DENOISING METHOD FOR ALEXNET CLASSIFICATION. THE PERFORMANCE ARE MULTI-CROP EVALUATION ON THE SAME 1000 IMAGES USED FOR DENOISING EXPERIMENTS.

Methods	σ				
	15	25	50	90	130
noise free	59.2 %	59.2 %	59.2 %	59.2 %	59.2 %
noisy	57.8 %	50.8 %	30.8 %	9.4 %	2.2 %
TRND	56.7 %	53.1 %	46.1 %	-	-
CBM3D	57.9 %	57.7 %	52.2 %	45.4 %	23.8 %
3DR+VGG-16	58.5 %	58.1 %	53.7 %	45.7 %	36.1 %
3DR+AlexNet	58.5 %	58.1 %	53.6 %	46.4 %	36.3 %

TABLE VI
COMPARISON WITH DIFFERENT DENOISING METHOD FOR VGG-16 CLASSIFICATION. THE PERFORMANCE ARE MULTI-CROP EVALUATION ON THE SAME 1000 IMAGES USED FOR DENOISING EXPERIMENTS.

Methods	σ				
	15	25	50	90	130
noise free	69.6 %	69.6 %	69.6 %	69.6 %	69.6 %
noisy	65.8 %	61.3 %	49.6 %	25.4 %	9.1 %
TRND	67.1 %	63.3 %	53.2 %	-	-
CBM3D	69.0 %	66.7 %	60.7 %	52.0 %	23.8 %
3DR+VGG-16	69.7 %	67.4 %	60.3 %	45.2 %	30.6 %
3DR+AlexNet	69.3 %	68.2 %	60.3 %	45.8 %	34.0 %

TABLE VII
COMPARISON WITH 3DR UNSUPERVISED DENOISING FOR ALEXNET AND VGG-16. THE PERFORMANCE ARE MULTI-CROP EVALUATION ON THE SAME 1000 IMAGES USED FOR DENOISING EXPERIMENTS.

Classification	σ		
	15	25	50
AlexNet	69.6 %	69.6 %	11.4 %
VGG-16	65.8 %	61.3 %	31.8 %

and VI. We conclude that denoising methods indeed help improving the classification accuracy. For example, in the case of $\sigma = 25$, our proposed method 3DR+AlexNet gains 7.3% and 6.9% advantage over the case of classification without applying denoising for both AlexNet and VGG-16. In the case of $\sigma = 50$, we outperform noisy image classification by 22.4% and 10.7% respectively. Tab. V and VI confirm our intuition that denoising step is indeed useful for classification. More importantly, if we can develop more powerful denoising method, it is more likely to improve the classification method by larger margin.

On the other hand, what if we let the classification model denoise the corrupted image? *i.e*. we use the softmax layer for classification as our loss function, at the meantime get rid of the denoising loss layer and freeze the weights of classification model. As it turns out (Tab. VII), denoising by classification performs way worse than applying common denoising method. For instance, AlexNet merely achieves 11.4% accuracy on classification task for $\sigma = 50$. The performance of VGG-16 also presents huge setback. Hence, it may not be a good idea to let classification model to make the image clean to itself, denoising model supervised by regression is still necessary.

Sometimes Noise helps. Besides, we find out that for certain cases AlexNet/VGG-16 fail to recognize neither the noise free nor denoised image. However, they can classify noisy images correctly. For instance, for AlexNet there exists 2.5% 3.5% 3.1% images under the condition $\sigma = 15, 25, 30$, for VGG-16 the rate is 1.4%, 1.8%, 1.7% respectively. Given there are 1000 categories for ImageNet, such rate of noise-improve-classify phenomenon cannot be simply considered as coincidence, on the contrary, it shows that AlexNet and VGG-16 are highly non-linear, quite sensitive to the noise. As shown in Fig. 14, AlexNet recognize the noisy soup bowl image with 19.22% confidence rate, but it cannot classify other images correctly. Fig. 15 demonstrate a similar example as well, noisy image is accurately classified to be Angora for VGG-16 with 30.06% confidence rate. Such counter-intuitive examples deserves more thorough study in later work, such that we can boost and stable neural network model.

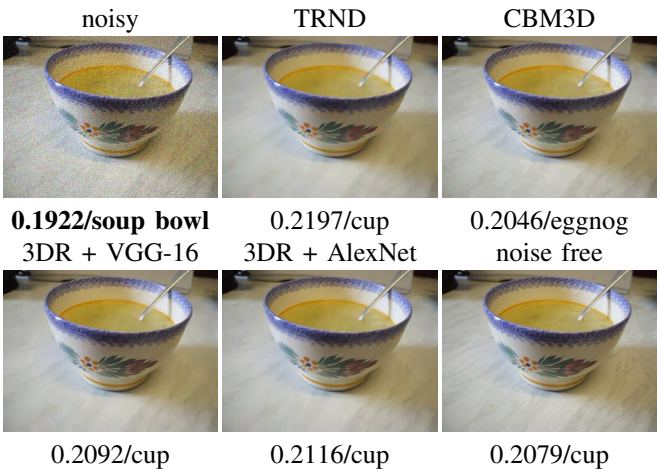

Fig. 14. AlexNet example for soup bowl, $\sigma = 25$.

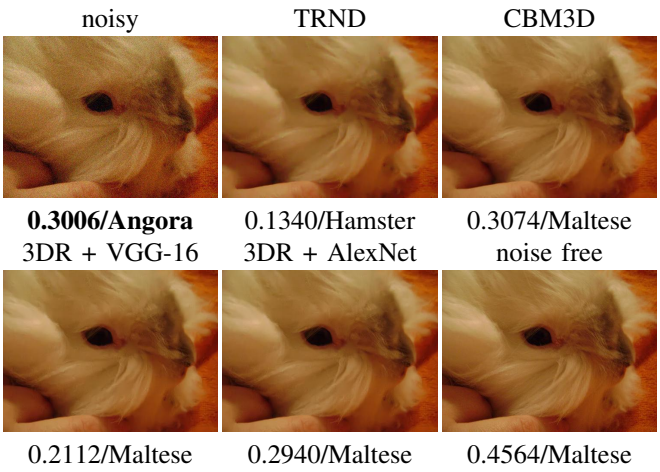

Fig. 15. VGG-16 example for Angora, $\sigma = 15$.

IV. CONCLUSION

In this paper, we proposed a novel color image denoising CNN model composed from a novel 3D residual learning stage and a standard top classification architecture (AlexNet/VGG-16). Experimental results on large and diverse datasets show that our model outperforms state-of-the-art methods. Meanwhile, we also study how denoising methods affect the classification models and conclude that denoising can indeed recover the classification accuracy. Last but not least, we notice that in certain cases noise can actually help classification models to 'see' the image clearer, which motivates us to study how to robustify neural networks by various noises in future work.

REFERENCES

[1] H. Brezis. *Functional analysis, Sobolev spaces and partial differential equations*. Springer Science & Business Media, 2010.
[2] H. Burger, C. Schuler, and S. Harmeling. Image denoising: Can plain neural networks compete with bm3d? In *IEEE Computer Vision and Pattern Recognition*, pages 2392–2399, 2012.
[3] Y. Chen, T. Pock, R. Ranftl, and H. Bischof. Revisiting loss-specific training of filter-based mrfs for image restoration. In *Pattern Recognition*, pages 271–281. Springer, 2013.

[4] Y. Chen, W. Yu, and T. Pock. On learning optimized reaction diffusion processes for effective image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5261–5269, 2015.

[5] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Color image denoising via sparse 3d collaborative filtering with grouping constraint in luminance-chrominance space. In *2007 IEEE International Conference on Image Processing*, volume 1, pages I–313. IEEE, 2007.

[6] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3d transform-domain collaborative filtering. *IEEE Trans. Image Processing*, 16:2080–2095, 2007.

[7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.

[8] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136, 2015.

[9] S. Gu, L. Zhang, W. Zuo, and X. Feng. Weighted nuclear norm minimization with application to image denoising. In *CVPR*, 2014.

[10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015.

[11] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(1):221–231, 2013.

[12] J. Kim, J. K. Lee, and K. M. Lee. Accurate image super-resolution using very deep convolutional networks. *arXiv preprint arXiv:1511.04587*, 2015.

[13] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.

[15] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

[16] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755. Springer, 2014.

[17] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman. Non-local sparse models for image restoration. In *IEEE 12th International Conference on Computer Vision*, pages 2272–2279, 2009.

[18] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 2, pages 416–423. IEEE, 2001.

[19] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 427–436. IEEE, 2015.

[20] A. Rajwade, A. Rangarajan, and A. Banerjee. Image denoising using the higher order singular value decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(4):849–862, 2013.

[21] U. Schmidt, J. Jancsary, S. Nowozin, S. Roth, and C. Rother. Cascades of regression tree fields for image restoration. 2014.

[22] U. Schmidt and S. Roth. Shrinkage fields for effective image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2774–2781, 2014.

[23] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

[24] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015.

[25] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

[26] R. Timofte, R. Rothe, and L. Van Gool. Seven ways to improve example-based single image super resolution. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

[27] R. Timofte, V. Smet, and L. Gool. Anchored neighborhood regression for fast example-based super-resolution. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1920–1927, 2013.

[28] R. Vemulapalli, O. Tuzel, and M.-Y. Liu. Deep gaussian conditional random field network: A model-based deep network for discriminative denoising. *arXiv preprint arXiv:1511.04067*, 2015.

[29] S. Zagoruyko and N. Komodakis. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

[30] R. Zeyde, M. Elad, and M. Protter. On single image scale-up using sparse-representations. In *Curves and Surfaces*, pages 711–730. Springer, 2010.

[31] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *Advances in neural information processing systems*, pages 487–495, 2014.

[32] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *IEEE International Conference on Computer Vision*, pages 479–486, 2011.

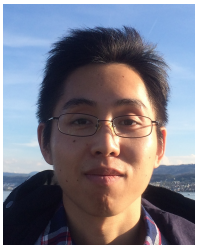

Jiqing Wu received the B.S. degree in mechanical engineering from Shanghai Maritime University, China, in 2006, the B.S. degree from TU Darmstadt, Germany, in 2012, the M.Sc. degree from ETH Zurich, Switzerland, in 2015, in mathematics, respectively. He is currently pursuing a PhD degree under the supervision of prof. Luc Van Gool, in his lab at ETH Zurich. His research interests mainly concern image demosaicing, image restoration, and early vision.

Radu Timofte obtained the PhD degree in Electrical Engineering at the KU Leuven, Belgium, in 2013, the M.Sc. at the Univ. of Eastern Finland in 2007 and the Dipl.-Ing. at the Technical Univ. of Iasi, Romania, in 2006. He worked as a researcher for the Univ. of Joensuu (2006-2007), the Czech Technical Univ. (2008) and the KU Leuven (2008-2013). Since 2013 he is postdoc at ETH Zurich, Switzerland, in the lab of prof. Luc Van Gool. He serves as a reviewer for major journals such as PAMI, TIP, IJCV, TNNLS, TKDE, PRL, and T-ITS, and conferences such as CVPR, ICCV, ECCV, and NIPS. He received best paper awards at ICPR 2012, at CVVT workshop from ECCV 2012 and at ChaLearn workshop from ICCV 2015. His current research interests include sparse and collaborative representations, image restoration and enhancement, and multi-view object class recognition.

Zhiwu Huang received the B.S. degree in computer science and technology from Huaqiao University, Quanzhou, China, in 2007, the M.S. degree in computer software and theory from Xiamen University, Xiamen, China, in 2010, and the Ph.D. degree in computer science and technology from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, in 2015. He has been a Post-Doctoral Researcher with the Computer Vision Laboratory, Swiss Federal Institute of Technology, Zurich, Switzerland, since 2015. His current research interests include Riemannian geometry, kernel learning, metric learning and feature learning with application to face recognition, face retrieval, facial expression recognition, gesture recognition, action recognition, and image set classification in computer vision.

Luc Van Gool obtained the master and PhD degrees in Electrical Engineering at the KU Leuven, Belgium, resp. in 1982 and 1991. He is a full professor for Computer Vision at both KU Leuven and ETH Zurich. With his two research labs, he focuses on object recognition, tracking and gesture analysis, and 3D acquisition and modeling. Luc Van Gool was a program chair of ICCV 2005, and general chair of ICCV 2011, and acted as GC of ECCV 2014. He is an editor-in-chief of Foundations and Trends in Computer Graphics and Vision. He also is a co-founder of the spin-off companies Eyetronics, GeoAutomation, kooaba, procedural, eSaturnus, upicto, Fashwell, Merantix, Spectando, and Parquery. He received several best paper awards, incl. at ICCV 1998, CVPR 2007, ACCV 2007, ICRA 2009, BMVC 2011, and ICPR 2012.